

APPROXIMATE MESSAGE PASSING UNDER FINITE ALPHABET CONSTRAINTS

Andreas Müller*

Dino Sejdinovic†

Robert Piechocki*

* Merchant Venturers School of Engineering, University of Bristol, Bristol, BS8 1UB, UK

† Gatsby Computational Neuroscience Unit, University College London, London, WC1N 3AR, UK
{andreas.muller, r.j.piechocki}@bristol.ac.uk, dino@gatsby.ucl.ac.uk

ABSTRACT

In this paper we consider *Basis Pursuit De-Noising* (BPDN) problems in which the sparse original signal is drawn from a finite alphabet. To solve this problem we propose an iterative message passing algorithm, which capitalises not only on the sparsity but by means of a prior distribution also on the discrete nature of the original signal. In our numerical experiments we test this algorithm in combination with a Rademacher measurement matrix and a measurement matrix derived from the random demodulator, which enables compressive sampling of analogue signals. Our results show in both cases significant performance gains over a linear programming based approach to the considered BPDN problem. We also compare the proposed algorithm to a similar message passing based algorithm without prior knowledge and observe an even larger performance improvement.

Index Terms— Compressive Sampling, Signal Recovery, Finite Alphabet, Message Passing

1. INTRODUCTION

When the information bearing part of a signal lies only in a small sub-space of the entire signal space, a uniform full rate sampling approach, i.e., sampling at Nyquist rate, is inefficient. In [1] and [2], Donoho and Candès *et al.* addressed this observation and introduced *Compressive Sampling* (CS). The key contribution of these papers was to show how random matrices and l_1 -minimization can be applied to achieve optimal recovery of a sparse signal from a very limited number of measurements.

While initial reconstruction algorithms were based on convex optimisation and linear programming, recent advances in compressive sampling have led to the development of various other algorithms which solve the reconstruction task at a lower computational complexity. Inspired by the success of message passing algorithms as used for the decoding operation of some channel codes, the authors of [3] solve the reconstruction of the original signal by means of belief propagation on a sparse graph. More recently, a simple iterative soft thresholding reconstruction algorithm, called *Approximate Message Passing* (AMP), has been proposed in

[4]. The basis of this algorithm is also belief propagation, albeit on a fully connected graph, and it exhibits a virtually equivalent sparsity-undersampling trade-off to that of linear programming based reconstruction algorithms.

In this paper, we address the problem of reconstructing a sparse finite alphabet signal from a limited number of noisy measurements which are obtained by compressive sampling. This problem appears in many areas such as spectrum sensing, symbol detection in digital communications, and multi-user detection, cf. [5, 6]. Many existing reconstruction algorithms for compressive sampling exploit the knowledge of the sparsity level of the original signal. Building upon the AMP framework, we propose here a novel AMP-based algorithm which capitalizes not only on the sparsity but also on a prior distribution, which manifests the finite alphabet property of the original signal.

In the derivation of this algorithm we assume that a time discrete sparse signal is applied directly to a measurement matrix whose entries are randomly sampled from $\{-\frac{1}{\sqrt{R}}, +\frac{1}{\sqrt{R}}\}$, where R is the number of measurements. In practice however, we are typically confronted with various types of measurement matrices. To this end we consider here as an application also the random demodulator [7], which allows to sample time-continuous analogue signals at sub-Nyquist rates and resorts to techniques from compressive sampling to reconstruct the signal. In both cases our numerical experiments indicate that the proposed algorithm offers an excellent performance in reconstructing sparse signals from noisy undersampled observations. When we compare the new algorithm with a standard linear programming based reconstruction algorithm and the AMP algorithm for the BPDN problem from [8], we observe a significant performance improvement. Moreover, we note that the computational cost of the algorithm proposed in this paper is comparable to existing state of the art algorithms.

The remainder of the paper is organized as follows. In Section 2 we outline the considered compressive sampling problem. Next, we introduce our proposed approximate message passing reconstruction algorithm with prior knowledge in Section 3. In Section 4.1 we compare the proposed algorithm at first with the linear programming based SPGL1

algorithm and the AMP algorithm for the BPDN problem in combination with a Rademacher measurement matrix. Then we briefly recap in Section 4.2 on the random demodulator and also present results for this scenario. Finally, Section 5 concludes the paper.

2. PROBLEM OUTLINE

Let $\mathcal{A} = \{a_1, \dots, a_S\}$ denote a finite set of S non-zero real numbers and define $\mathcal{A}_0 := \mathcal{A} \cup \{0\}$. The W -dimensional sparse column vector \mathbf{b} shall have only $K \ll W$ non-zero entries, which are drawn uniformly at random from \mathcal{A} . Moreover, let $\Psi \in \mathbb{R}^{R \times W}$ be a matrix, for which $R \leq W$ holds. Given the noisy observation

$$\mathbf{v} = \Psi \mathbf{b} + \mathbf{n}, \quad (1)$$

where $\mathbf{n} = [n_0, \dots, n_r, \dots, n_{R-1}]^T$ and n_r is i.i.d. $\mathcal{N}(0, \sigma^2)$, our aim is to recover \mathbf{b} . Problems of this type are frequently considered in the CS literature. One way to reconstruct the sparse vector \mathbf{b} is by solving the optimization problem

$$\arg \min_{\hat{\mathbf{b}} \in \mathcal{A}_0^W} \|\hat{\mathbf{b}}\|_0 \quad \text{subject to} \quad \|\Psi \hat{\mathbf{b}} - \mathbf{v}\|_2^2 \leq \gamma', \quad (2)$$

for an optimization constant γ' , which is chosen depending on the noise variance. In general though the l_0 -minimization used in (2) is NP-hard. This has led to various approximate algorithms, including those based on convex relaxation, i.e., on l_1 -minimization, like BPDN [9] or LASSO [10]. Many of the approximate algorithms for CS can thus be applied directly to approximate the solution of (2) by solving the relaxed problem

$$\arg \min_{\hat{\mathbf{b}} \in \mathbb{R}^W} \|\hat{\mathbf{b}}\|_1 \quad \text{subject to} \quad \|\Psi \hat{\mathbf{b}} - \mathbf{v}\|_2^2 \leq \gamma, \quad (3)$$

and ensuring by, for example, thresholding that $\hat{\mathbf{b}} \in \mathcal{A}_0^W$ holds. However, in addition to the number of non-zero entries in the sparse W -dimensional signal \mathbf{b} , the knowledge that its entries lie in the set \mathcal{A}_0 can also be utilized directly to improve the recovery of \mathbf{b} , as will be discussed in the following sections.

3. AMP WITH DISCRETE PRIOR DISTRIBUTION

It is well known that the solution of the l_1 -minimization in (3) corresponds to a mode of the posterior distribution when a double-exponential prior distribution is used. Donoho *et al.* also use a double-exponential prior in the derivation of their AMP algorithm [4], where the contribution of the prior distribution in the message update rules of the belief propagation algorithm can be interpreted as a ‘sparsity promoting’ soft thresholding operation. However, their results are quite general and a similar approach can be applied when a different choice of prior distribution is more suitable. Indeed, [8]

argues that an estimate of the input distribution can be used to improve the recovery algorithm. Therefore, it is natural to employ a discrete prior distribution when the original signal is drawn from a finite alphabet, and we take

$$f(b_w) = \pi_0 \cdot \delta\{b_w = 0\} + \sum_{s=1}^S \pi_s \cdot \delta\{b_w = a_s\} \quad (4)$$

as prior for each b_w in $\mathbf{b} = [b_0, \dots, b_w, \dots, b_{W-1}]^T$, where $\pi_0 = 1 - \frac{K}{W}$, and $\pi_s = \frac{K}{WS}$ for $1 \leq s \leq S$. Note that this prior is constructed under the assumption that the non-zero entries in \mathbf{b} are drawn uniformly from \mathcal{A} . Should some non-zero entries be more likely than others, the prior distribution can be easily modified to reflect this additional information.

Consider a fully connected bipartite graph between R measurement nodes on one side and W variable nodes on the other side. The measurement nodes shall represent the entries in \mathbf{v} and likewise the variable nodes the unknown entries in \mathbf{b} . As in [4] we study belief propagation (BP) message updates between the measurement and variable nodes on this complete graph. We denote the set of the W variable nodes as $[W]$ and use $w, \omega \in \{0, 1, \dots, W-1\}$ as indices for this set. Similarly, we apply $r, \rho \in \{0, 1, \dots, R-1\}$ as indices for the set $[R]$ of all R measurement nodes. At iteration t we denote the message passed from the variable node w to the measurement node r by $\nu_{w \rightarrow r}^{(t)}(b_w)$ and the message passed on this edge in the opposite direction by $\hat{\nu}_{r \rightarrow w}^{(t)}(b_w)$. In the t -th iteration the BP message updates are then given as

$$\begin{aligned} \hat{\nu}_{r \rightarrow w}^{(t)}(b_w) &\propto \int \exp \left[-\frac{1}{2\sigma^2} (v_r - (\Psi \mathbf{b})_r)^2 \right] \\ &\quad \cdot \prod_{\omega \neq w} \nu_{\omega \rightarrow r}^{(t-1)}(b_\omega) d\mathbf{b}_{-w}, \end{aligned} \quad (5)$$

$$\nu_{w \rightarrow r}^{(t)}(b_w) \propto f(b_w) \prod_{\rho \neq r} \hat{\nu}_{\rho \rightarrow w}^{(t)}(b_w), \quad (6)$$

where $d\mathbf{b}_{-w}$ denotes that integration is over all variables except b_w , and σ^2 is the variance of the noise in the measurements. Note that the messages from variable to measurement nodes in (6) are proportional to the probability mass function on \mathcal{A}_0 . We initialize

$$\nu_{w \rightarrow r}^{(0)}(b_w) = f(b_w). \quad (7)$$

We denote the mean and the variance of the message in (6) by $\xi_{w \rightarrow r}^{(t)}$ and $\tau_{w \rightarrow r}^{(t)}$, respectively. Note that $\xi_{w \rightarrow r}^{(0)}$ and $\tau_{w \rightarrow r}^{(0)}$ are initialized to the prior mean and variance $\forall w \in [W], \forall r \in [R]$. We will derive the algorithm here under the assumption that $\psi_{r,w} \in \left\{ -\frac{1}{\sqrt{R}}, +\frac{1}{\sqrt{R}} \right\}$. Later in our simulations however we relax this assumption and work with a general Ψ with normalized columns.

Consider the random vector $\mathbf{b}_{-w} = [b_0, b_1, \dots, b_{w-1}, b_{w+1}, \dots, b_{W-1}]^T$ distributed according to the product measure $\prod_{\omega \neq w} \nu_{\omega \rightarrow r}^{(t)}(b_\omega)$, and the associated scalar random




Fig. 1: Detection error rate $P(\hat{b}_w \neq b_w)$ versus the noise variance σ^2 when the measurement matrix Ψ in (1) is a Rademacher matrix with unit column norm, or the random demodulator matrix. The number of iterations for both AMP algorithms is set to $T = 50$.

variable

$$x_{r \rightarrow w}^{(t)} = v_r - \sum_{\omega \neq w} \psi_{r,\omega} b_\omega. \quad (8)$$

Denote the induced density of $x_{r \rightarrow w}^{(t)}$ by g . By the central limit theorem, for large W , g can be approximated by a Gaussian density with mean $v_r - \sum_{\omega \neq w} \psi_{r,\omega} \xi_{\omega \rightarrow r}^{(t)}$ and variance $\frac{1}{R} \sum_{\omega \neq w} \tau_{\omega \rightarrow r}^{(t)}$. Thereby, if we write (5) in its equivalent form

$$\hat{\nu}_{r \rightarrow w}^{(t)}(b_w) \propto \mathbb{E}_{x_{r \rightarrow w}^{(t)} \sim g} \exp \left[-\frac{1}{2\sigma^2} (x_{r \rightarrow w}^{(t)} - \psi_{r,w} b_w)^2 \right],$$

we see that the factor to variable message update (5) can be approximated by a Gaussian integral. These observations lead to the simplified algorithm for message passing with a discrete prior distribution given in Algorithm 1¹. Note that we track the posterior probabilities in the log-domain, as numerical simulations indicate that this results in a numerically more stable implementation of the simplified message-passing algorithm.

4. NUMERICAL EXPERIMENTS

4.1. Rademacher Measurement Matrix

For the numerical experiments presented here we assume that $\mathcal{A}_0 = \{-1, 0, +1\}$ and fix $W = 512$, $R = 205$, $K = 20$. The entries in the measurement matrix Ψ shall be drawn uniformly at random from $\{-1/\sqrt{R}, +1/\sqrt{R}\}$, which makes

¹Throughout the algorithm, “dot” is a placeholder for any $r \in [R]$ or for ρ , in which case the summation in Step (3) is over all $\rho \in [R]$.

Ψ a Rademacher matrix with unit column norm. To obtain the estimate $\hat{\mathbf{b}}$ for the original \mathbf{b} from the noisy observation \mathbf{v} we apply the approximate message passing algorithm, as outlined in Algorithm 1. For a comparison we also apply to the same problem the algorithms for the BPDN problem from [11, 12], named SPGL1-BPDN, and from [8], named AMP-BPDN. Unlike the proposed AMP algorithm the SPGL1-BPDN and the AMP-BPDN algorithm require the parameter γ as input, which we choose as $\gamma = \sqrt{\lceil W/R \rceil} \cdot \sigma \cdot \sqrt{R} \cdot \sqrt{1 + \sqrt{2}/\sqrt{R}}$. In contrast to our proposed algorithm, which always returns an estimate for \mathbf{b} in $\{-1, 0, +1\}^W$, the SPGL1-BPDN and the AMP-BPDN algorithm return estimates in \mathbb{R}^W . For this reason one can threshold the $\hat{\mathbf{b}}$ obtained from these two algorithms by α and take as the output $\text{sign}(\hat{b}_w)$, where $|\hat{b}_w| \geq \alpha$, and zero otherwise. Alternatively, one can search for the K entries in $\hat{\mathbf{b}}$ with the largest magnitude and set them depending on their sign to ± 1 and all other entries to zero.

In Figure 1 the detection error rate, i.e., $P(\hat{b}_w \neq b_w)$, is plotted for AMP-BPDN, SPGL1-BPDN, and the algorithm proposed in this paper versus the noise variance. Our simulations indicate that the detection error rate of the SPGL1-BPDN and AMP-BPDN algorithm are strongly dependent on the chosen threshold α and that the decoding rule which simply chooses the K largest entries of $\hat{\mathbf{b}}$ achieves the best performance for both of these algorithms (this version is plotted in Figure 1). However, even in this case the detection error probability performance of the SPGL1-BPDN and the AMP-BPDN algorithm is approximately 6dB respectively 20dB worse than that of the proposed AMP algorithm with discrete prior.

4.2. Random Demodulator

In this section we consider discrete multi-tone signals, which occur, for example, in orthogonal frequency division multiplex systems, in combination with the random demodulator. For this class, it is shown in [7] how the operation of the random demodulator, which is depicted in Figure 2, on the analogue signal can be described equivalently by a time-discrete representation.

Let $\mathbf{F} \in \mathbb{C}^{W \times W}$ be a discrete Fourier transform matrix and $\mathbf{b} \in \{-1, 0, +1\}^W$ a K -sparse vector. A time discrete representation of the analogue multi-tone signal is then given




Fig. 2: Structure of the random demodulator as discussed in [7]

by $\mathbf{x} = \mathbf{F}\mathbf{b}$. The multiplication operation of the signal \mathbf{x} with the chipping sequence and the integrate and dump operation of the random demodulator are described by the $W \times W$ diagonal matrix \mathbf{D} and by $\mathbf{H} \in \{0, 1\}^{R' \times W}$, respectively. \mathbf{H} shall have W/R' consecutive ones in the r -th row starting from column $rW/R' + 1$, where $r = 0, \dots, R' - 1$. As proposed in [7] we allow \mathbf{H} to have fractional elements in some of its columns when R does not divide W . To summarize, the sparse signal \mathbf{b} is observed by the random demodulator through the random demodulator matrix $\Psi' = \mathbf{H}\mathbf{D}\mathbf{F} \in \mathbb{C}^{R' \times W}$. The sign of the non-zero entries of \mathbf{D} on the main diagonal is chosen independently at random. However, we choose the magnitude of these entries here such that the columns of Ψ' have unit norm.

As in Section 4.1 we aim here to recover \mathbf{b} from the noisy samples \mathbf{v} given in (1), where however the real valued measurement matrix

$$\Psi = \begin{bmatrix} \text{Re}\{\Psi'\} \\ \text{Im}\{\Psi'\} \end{bmatrix} \quad (9)$$

depends now on the random demodulator matrix Ψ' . For our numerical experiments with the measurement matrix given by (9), we set $W = 512$, $R = 204$, $R' = 102$, and $K = 20$. For this scenario the detection error rate, i.e., $P(\hat{b}_w \neq b_w)$, is also plotted in Figure 1 versus the noise variance. The performance of all three algorithms is very similar to the case in Section 4.1 with a Rademacher measurement matrix. The performance gap between the SPGL1-BPDN, the AMP-BPDN, and the proposed AMP algorithm with discrete prior is here also approximately 6dB and 20dB wide.

5. CONCLUSIONS

In this paper we have developed a novel reconstruction algorithm for compressive sampling problems, which applies the prior knowledge that the entries of the original signal vector belong to a finite alphabet. Our simulation results show for this algorithm significant performance gains over existing reconstruction algorithms for the BPDN problem. In future work we hope to further simplify the message passing update equations of the proposed algorithm and thus further reduce its complexity.

References

- [1] D.L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [2] E.J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [3] D. Baron, S. Sarvotham, and R.G. Baraniuk, "Bayesian compressive sensing via belief propagation," *IEEE Transactions on Signal Processing*, vol. 58, no. 1, pp. 269–280, 2010.
- [4] D.L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. of National Academy of Sciences (PNAS)*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.

Algorithm 1 AMP with discrete prior

- **Input:** Measurement matrix Ψ , observation \mathbf{v} , noise variance σ^2 , alphabet \mathcal{A}_0 , prior probabilities $\pi_0, \pi_1, \dots, \pi_S$, number of iterations T

- **Output:** signal estimate $\hat{\mathbf{b}}$

1. Initialize: $t = 1$, and $l_{w \rightarrow r}^{(0)}(a_s) = \pi_s$, for $w \in [W]$, $r \in [R]$, $s \in [S]$.
2. Calculate the mean and variance of the variable-to-factor messages:

$$\xi_{w \rightarrow \cdot}^{(t)} = \sum_{s=0}^S a_s \exp l_{w \rightarrow \cdot}^{(t-1)}(a_s)$$

$$\tau_{w \rightarrow \cdot}^{(t)} = \sum_{s=0}^S a_s^2 \exp l_{w \rightarrow \cdot}^{(t-1)}(a_s) - \left(\xi_{w \rightarrow \cdot}^{(t)}\right)^2$$

3. Approximate the mean and variance of $\prod_{\rho \neq \cdot} \hat{\nu}_{\rho \rightarrow w}^{(t)}(b_w)$:

$$\mu_{w \rightarrow \cdot}^{(t)} = \sum_{\rho \neq \cdot} \psi_{\rho, w} \left(v_{\rho} - \sum_{\omega \neq w} \psi_{\rho, \omega} \xi_{\omega \rightarrow \rho}^{(t)} \right)$$

$$\eta_{w \rightarrow \cdot}^{(t)} = \frac{1}{R} \sum_{\omega \neq w} \tau_{\omega \rightarrow \cdot}^{(t)} + \sigma^2$$

4. Incorporate prior and normalize:

$$\bar{l}_{w \rightarrow \cdot}^{(t)}(a_s) = \log \pi_s - \frac{\left(\mu_{w \rightarrow \cdot}^{(t)} - a_s\right)^2}{2\eta_{w \rightarrow \cdot}^{(t)}}$$

$$l_{w \rightarrow \cdot}^{(t)}(a_s) = -\log \left(\sum_{s'=0}^S \exp \bar{l}_{w \rightarrow \cdot}^{(t)}(a_{s'}) \right) + \bar{l}_{w \rightarrow \cdot}^{(t)}(a_s)$$

Set $t \leftarrow t + 1$, and repeat (2)-(4) until stopping criterion holds.

5. Return $\hat{\mathbf{b}} = [\hat{b}_0, \dots, \hat{b}_{W-1}]^T$, where

$$\hat{b}_w = \arg \max_{a \in \mathcal{A}_0} l_{w \rightarrow \cdot}^{(t)}(a).$$

- [5] Hao Zhu and G.B. Giannakis, "Sparsity-embracing multiuser detection for cdma systems with low activity factory," in *IEEE International Symposium on Information Theory (ISIT)*, Jul. 2009, pp. 164–168.
- [6] Zhi Tian, G. Leus, and V. Lottici, "Detection of sparse signals under finite-alphabet constraints," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2009, pp. 2349–2352.
- [7] J.A. Tropp, J.N. Laska, M.F. Duarte, J.K. Romberg, and R.G. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 520–544, Jan. 2010.
- [8] D.L. Donoho, A. Maleki, and A. Montanari, "Message passing algorithms for compressed sensing: I. motivation and construction," in *IEEE Information Theory Workshop (ITW)*, Jan. 2010, pp. 1–5.
- [9] S. Chen, D.L. Donoho, and M.A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, 2001.
- [10] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [11] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction," June 2007, <http://www.cs.ubc.ca/labs/scl/spgl1>.
- [12] E. van den Berg and M. P. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.